

0020201-86521250

Thus, a method and apparatus for dynamically replicating selected resources in computer networks is provided. One skilled in the art will appreciate that the present invention can be practiced by other than the described embodiments, which are presented for purposes of illustration and not limitation, and the present invention is limited only by the claims that follow.

1 What is claimed:

2       1. A method of processing resource requests in a computer network, the  
3       method comprising,

4       (i) by a client:

5           (A) making a request for a particular resource from an origin server,  
6           the request including a resource identifier for the particular  
7           resource;

8       (ii) by a reflector:

9           (B) intercepting the request from the client to the origin server;  
10          (C) selecting a repeater to process the request;  
11          (D) providing to the client a modified resource identifier designating  
12           the repeater;

13       (iii) by the client:

14           (E) receiving the modified resource identifier from the reflector; and  
15           (F) making a request for the particular resource from the repeater  
16           designated in the modified resource identifier;

17       (iv) by the repeater:

18           (G) receiving the request from the client; and  
19           (H) ~~returning the requested resource to the client.~~ 2. A method

20       as in claim 1 further comprising, by the repeater:

21           (I) making a request for the resource from the origin server; and  
22           (J) receiving the resource from the origin server.

23

24           3. A method as in claim 1 wherein the selecting of a repeater by the  
25 reflector comprises:

26           (C1) partitioning the network into groups;  
27           (C2) determining which group the client is in;  
28           (C3) selecting, from a plurality of repeaters in the network, a set of repeaters  
29           having a lowest cost relative to the group which the client is in; and  
30           (C4) selecting as the repeater a member of the selected set of repeaters.

31

32           4. A method as in claim 3, wherein the cost of a repeater is a value based on  
33 that repeater's current load and a maximum load for that repeater.

34

35           5. A method as in claim 3, wherein the cost of a repeater is a value based on  
36 a predicted cost or speed of transmission between the repeater and a client in the group.

37

38           6. A method as in claim 1 wherein the particular resource itself contains at  
39 least one other resource identifier of at least one other resource, the method further  
40 comprising:

41           rewriting the particular resource to replace at least some of the resource  
42 identifiers contained therein with modified resource identifiers designating a repeater  
43 instead of the origin server.

44

45           7. A method as in claim 6 wherein the rewriting is performed by one of the  
46 repeater, the reflector or another repeater.

47

48           8. A method of processing resource requests in a computer network, the  
49 method comprising,

50           (i) by a client:

(A) making a request for a particular resource from an origin server, the request including a resource identifier for the particular resource;

(ii) by a reflector:

(B) intercepting the request from the client to the origin server;

(C) determining whether to reflect the request to a repeater;

(D) when the reflector determines not to reflect the request, forwarding the request to the origin server, otherwise

(D1) selecting a repeater to process the request;

(D2) providing to the client a modified resource identifier designating the repeater.

9. A method as in claim 8, further comprising, when the reflector determines to reflect the request,

(iii) by the client:

(E) receiving the modified resource identifier from the reflector; and

(F) making a request for the particular resource from the repeater designated in the modified resource identifier;

(iv) by the repeater:

(G) receiving the request from the client; and

(H) returning the requested resource to the client.

10. A method as in claim 8 wherein the reflector determines whether to reflect a request by comparing the resource identifier with regular expression patterns of repeatable resources.

77        11. A method as in claim 8; wherein the reflector has a threshold aggregate  
78 information rate (TAIR) associated therewith, and wherein the determining of whether  
79 to reflect the request to a repeater comprises:

80            determining whether the TAIR of the reflector is exceeded by a measured  
81 aggregate information rate (MAIR) for the reflector, wherein the reflector determines  
82 not to reflect the request when the MAIR exceeds the TAIR for the reflector.

83

84        12. A method as in claim 8, wherein the reflector has a threshold aggregate  
85 information rate (TAIR) associated therewith, and wherein the determining of whether  
86 to reflect the request to a repeater comprises:

87            probabilistically determining whether the TAIR of the reflector is exceeded by a  
88 measured aggregate information rate (MAIR) for the reflector, wherein the reflector  
89 determines not to reflect the request as an exponential function of the difference  
90 between the MAIR and the TAIR.

91

92        13. A method as in any of claims 11-12, wherein the MAIR is obtained from  
93 repeaters according to the rate at which they have transmitted data on behalf of the  
94 reflector during a given time interval.

95

96        14. A method as in any one of claims 1-12 wherein the network is the  
97 Internet and wherein the resource identifier is a uniform resource locator (URL) for  
98 designating resources on the Internet, and wherein the modified resource identifier is a  
99 URL designating the repeater and indicating the reflector or origin server, and wherein  
100 the modified resource identifier is provided to the client using a REDIRECT message.

101

102        15. In a computer network wherein clients request resources from origin  
103 servers, a method comprising:

104            providing at least one repeater;

105 providing reflectors at some of the origin servers, each reflector intercepting  
106 client resource requests made to its respective origin server; and  
107 each reflector selectively redirecting client resource requests for certain resources  
108 to one of the repeaters.

109

110 16. A method as in claim 15 further comprising, by repeaters in the network:  
111 servicing redirected client resource requests; and  
112 selectively maintaining copies of requested resources,  
113 whereby resources corresponding to redirected resource requests are selectively  
114 migrated from their origin servers to one or more repeaters.

115

116 17. A computer network comprising:  
117 a plurality of origin servers, at least some of the origin servers having reflectors  
118 associated therewith;  
119 a plurality of repeaters; and  
120 a plurality of clients,  
121 wherein each reflector is adapted to intercept resource requests made to its  
122 respective origin server and to selectively redirect the resource requests to a dynamically  
123 selected repeater.

124

125 18. In a computer network wherein clients request resources from origin  
126 servers, a reflector mechanism associated with an origin server, the reflector mechanism  
127 comprising:  
128 means for intercepting a resource request made by client of an origin server;  
129 means for analyzing the resource request to determine whether to service the  
130 request locally at the origin server;  
131 means for determining a best repeater in the network to service the request when

132 the analyzing means determines that the request should not be serviced locally; and  
133 means for redirecting the client to the best repeater.

134

135 19. A reflector mechanism as in claim 18 wherein the network is partitioned  
136 into groups and the means for determining the best repeater comprises:  
137 means for determining which group the client is in;  
138 means for selecting, from a plurality of repeaters in the network, a set of  
139 repeaters having a lowest cost relative to the group the client is in; and  
140 means for selecting as the best repeater a member of the set of repeaters.

141

142 20. A reflector mechanism as in claim 19, wherein the cost of a repeater is a  
143 value based on a predicted cost or speed of transmission between the repeater and a  
144 client in the group.

145

146 21. A mechanism as in claim 19, wherein the cost of a repeater is a value  
147 based on that repeaters current load and a maximum load for that repeater.

148

149 22. A reflector as in claim 16 wherein the resource itself contains resource  
150 identifiers, the reflector further comprising:  
151 means for rewriting the resource to replace at least some of the resource  
152 identifiers contained therein with modified resource identifiers designating the repeater  
153 instead of the origin server.

154

155 23. In a computer network wherein clients request resources from origin  
156 servers, a repeater mechanism comprising:  
157 means for receiving a resource request from a client;  
158 means for determining whether the resource is available locally;  
159 means for, when it is determined that the resource is not available locally,

obtaining the resource from an origin server; and  
means for providing the resource to the client.

24. A reflector as in claim 18 wherein the resource itself contains resource  
5 identifiers, the repeater further comprising:  
means for rewriting the resource to replace at least some of the resource  
identifiers contained therein with modified resource identifiers designating the repeater  
instead of the origin server.

add a1>  
add c1>  
add d1>  
add f8>

00000000000000000000000000000000